

Ceph-deploy - Bug #12557

systemd: osd create should enable ceph.target

07/31/2015 07:23 PM - Sage Weil

Status: Resolved	% Done: 0%
Priority: Urgent	
Assignee:	
Category:	
Target version: 1.5.29	
Source: other	Reviewed:
Tags:	Affected Versions:
Backport:	ceph-qa-suite:
Regression: No	Crash signature (v1):
Severity: 3 - minor	Crash signature (v2):
Description	
Related issues:	
Related to Ceph-deploy - Bug #12726: ceph-deploy suite fails on CentOS 7	Closed 08/18/2015

Associated revisions

Revision 9cfe20a9 - 09/08/2015 08:28 PM - Sage Weil

osd: enable ceph on osd create

Fixes: #12557

Signed-off-by: Sage Weil <sage@redhat.com>

History

#1 - 08/26/2015 09:52 PM - Travis Rhoden

Is this still an issue? I believe that it might be. I'm looking at ceph-deploy's code for osd prepare/activate/create, and I have to admit there are parts of it that have never made sense to me.

Namely, i think of "create" as "activate + prepare". Except that is **not** what ceph-deploy does. In ceph-deploy, "ceph-deploy osd prepare" runs "ceph-disk prepare" on the remote node and "ceph-deploy osd activate" runs "ceph-disk activate", but "ceph-deploy osd create" runs "ceph-disk prepare" followed by "udevadm settle". I have no idea why it is implemented this way. instead of running "ceph-disk activate" after "ceph-disk prepare".

The code in ceph-deploy to enable ceph.target currently is only in the activate() function, which only gets triggered if you call ceph-deploy osd activate. This means that if you use ceph osd create, it's not going to get called. That's probably what is causing some of the other issues we have seen with OSDs not starting under systemd... Just a guess.

#2 - 09/04/2015 04:14 PM - Sage Weil

Travis Rhoden wrote:

Is this still an issue? I believe that it might be. I'm looking at ceph-deploy's code for osd prepare/activate/create, and I have to admit there are parts of it that have never made sense to me.

Namely, i think of "create" as "activate + prepare". Except that is **not** what ceph-deploy does. In ceph-deploy, "ceph-deploy osd prepare" runs "ceph-disk prepare" on the remote node and "ceph-deploy osd activate" runs "ceph-disk activate", but "ceph-deploy osd create" runs "ceph-disk prepare" followed by "udevadm settle". I have no idea why it is implemented this way. instead of running "ceph-disk activate" after "ceph-disk prepare".

It's because normally udev will see the created osd volume and do the activate step. We can change it to an explicit activate too, though.. that will work just as well (and will work properly when udev is broken, or when the do-no-activate flag is set).

```
> The code in ceph-deploy to enable ceph.target currently is only in the activate() function, which only gets triggered if you call ceph-deploy osd activate. This means that if you use ceph osd create, it's not going to get called. That's probably what is causing some of the other issues we have seen with OSDs not starting under systemd... Just a guess.
```

Ah.. yeah, that should be fixed. I think swapping udevadm settle for create will kill both birds with one stone?

#3 - 09/08/2015 08:30 PM - Sage Weil

- Status changed from New to Fix Under Review

#4 - 09/14/2015 06:24 PM - Travis Rhoden

- Status changed from Fix Under Review to Resolved

- Target version set to 1.5.29